

CLAIMS:

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1. A sampling-based system for adaptively optimizing a computer program executing in an execution environment, said execution environment including one or more compiler devices for providing various levels of program optimization, said system comprising:

a runtime measurements sub-system for monitoring execution of said computer program to be optimized, said monitoring including obtaining raw profile data samples and characterizing said raw profile data;

a controller device for receiving said characterized raw profile data from said runtime measurements sub-system and analyzing said data for determining whether a level of program optimization for said executing program is to be performed by a compiler device, said controller generating a compilation plan in accordance with a determined level of optimization; and,

a recompilation sub-system for receiving a compilation plan from said controller and invoking a compiler device for performing said level of program optimization of said executing program in accordance with said compilation plan.

2. The system as claimed in Claim 1, wherein said runtime measurements sub-system comprises one or more organizer devices for processing said raw profile data and characterizing said data as meeting a hotness threshold of activity.

3. The system as claimed in Claim 2, wherein said runtime measurements sub-system comprises:

a mechanism for counting a number of samples that are taken from the executing

4 program; and,

5 a mechanism for comparing the number of samples to a predetermined sampling
6 size threshold, and in response to the number of samples exceeding said sampling threshold,
7 invoking said organizer device to process said raw profile data.

1 4. The system as claimed in Claim 3, wherein said raw profile data samples relate to one or more
2 method activations in said executing program.

1 5. The system as claimed in Claim 4, wherein said organizer device comprises a mechanism for
2 comparing said raw profile data of method activations against a corresponding activity hotness
3 threshold for one or more methods, and identifying one or more methods as meeting said activity
4 hotness threshold for input to said controller device.

1 6. The system as claimed in Claim 3, wherein said controller device adaptively adjusts said
2 sampling size threshold.

1 7. The system as claimed in Claim 4, wherein a level of program optimization includes
2 recompiling an executing method, said controller device further comprising mechanism for
3 adapting said sampling size threshold in accordance with an amount of recompilation that occurs.

1 8. The system as claimed in Claim 2, wherein said controller device dynamically adjusts said
2 activity hotness threshold to adapt to a current behavior of the executing computer program.

1 9. The system as claimed in Claim 4, wherein said controller instructs the recompilation
2 subsystem to insert intrusive profiling for one or more identified methods.

1 10. The system as claimed in Claim 4, wherein a compilation plan generated by said controller
2 device comprises an identifier of a method to be optimized; and, an optimization level indicating
3 a degree of optimization to be applied for said identified method.

1 11. The system as claimed in Claim 4, wherein a level of program optimization includes
2 recompiling an executing method, said controller device including a mechanism for identifying a
3 recompilation level that minimizes expected future running time of a recompiled program.

1 12. The system as claimed in Claim 11, wherein said mechanism for identifying a recompilation
2 level includes:

3 mechanism for determining an expected time T_i the program will spend executing
4 a method "m" if said method is not recompiled;

5 mechanism for determining a cost C_j of recompiling said method at an
6 optimization level "j", for $i \leq j \leq N$; and,

7 mechanism for determining an expected time T_j the program will spend executing
8 said method in the future, if said method is recompiled at level "j"; and,

9 comparison mechanism for evaluating the expression $C_j + T_j < T_i$, whereby said
10 controller device decides to one of: generate compilation plan for directing
11 recompilation of "m" at level "j" if said expression is true, and, not recompile if
12 said expression is false.

1 13. The system as claimed in Claim 1, wherein said raw profile data samples are taken at method
2 prologue and back edge yield points.

1 14. The system as claimed in Claim 1, wherein said controller instructs the recompilation
2 subsystem to perform online feedback-directed optimizations based on feedback from the current
3 executing program.

1 15. The system as claimed in Claim 14, wherein said raw profile data samples relate to call
2 context information associated with methods called by said program, said feedback comprising
3 said call context information.

1 16. The system as claimed in Claim 14, wherein said raw profile data samples relate to current
2 program variable values, said feedback comprising a subset of values assigned to said variables
3 during program execution.

1 17. The system as claimed in Claim 14, wherein said raw profile data samples relate to control
2 flow execution within a method, said feedback comprising execution frequency of control flow
3 paths within said executing method.

1 18. The system as claimed in Claim 1, wherein said execution environment includes an
2 interpreter device.

1 19. A method for adaptively optimizing a computer program executing in an execution
2 environment, said execution environment comprising one or more compiler devices for providing
3 various levels of program optimization, said method comprising:

- 4 a) sampling said executing computer program to obtain raw profile data samples;
5 b) characterizing said raw profile data as meeting a threshold criteria;
6 c) analyzing said characterized raw profile data for determining whether a level of
7 program optimization for said executing program is to be performed by a compiler device, and
8 generating a compilation plan in accordance with a determined level of optimization; and,
9 d) when optimization is to be performed, invoking a compiler device for
10 optimizing said executing program in accordance with said compilation plan.

1 20. The method as claimed in Claim 19, wherein said characterizing step b) comprises:
2 processing said raw profile data to determine whether said data meets an activity hotness

3 ~~threshold.~~

1 21. The method as claimed in Claim 20, wherein said sampling includes the steps of:
2 counting said samples that are taken from the executing program; and,
3 comparing the amount of samples taken to a predetermined sampling size
4 threshold, whereby in response to an amount of samples exceeding said sampling threshold,
5 performing said characterizing step.

1 22. The method as claimed in Claim 21, wherein said raw profile data sampled relates to one or
2 more method activations in said executing program.

3 23. The method as claimed in Claim 22, wherein said step of processing said raw profile data
4 comprises the steps of:
5 comparing said raw profile data of method activations against said corresponding
6 activity hotness threshold for one or more methods; and
7 identifying one or more methods as meeting said activity hotness threshold.

8 24. The method as claimed in Claim 21, further including the step of adaptively adjusting said
9 sampling size threshold.

1 25. The method as claimed in Claim 22, wherein said step of optimizing includes recompiling an
2 executing method meeting an activity hotness threshold, said method further including the step
3 of adapting said sampling size threshold in accordance with an amount of recompilation that
4 occurs.

1 26. The method as claimed in Claim 22, further including the step of dynamically adjusting said
2 activity hotness threshold for adapting to a current behavior of the executing computer program.

1 27. The method as claimed in Claim 26, wherein said step of dynamically adjusting said activity
 2 hotness threshold further includes the steps of:
 3 determining an amount of methods characterized as meeting said threshold criteria
 4 after one or more sampling periods;
 5 comparing said amount to a limit; and,
 6 in response to said comparing, one of: decreasing said activity hotness threshold if
 7 said limit is not met, and increasing the threshold if said limit is met or exceeded in said one or
 8 more sampling periods.

1 28. The method as claimed in Claim 22, further including the step of inserting intrusive profiling
 2 for one or more identified methods.

1 29. The method as claimed in Claim 24, wherein said step of generating a compilation plan
 2 includes: providing an identifier of said method to be optimized; and, an optimization level
 3 indicating a degree of optimization to be applied for said identified method.

1 30. The method as claimed in Claim 22, wherein said analyzing step c) further includes
 2 identifying a recompilation level that minimizes expected future running time of a recompiled
 3 version.

1 31. The method as claimed in Claim 30, wherein said step of identifying a recompilation level
 2 includes the steps of:
 3 determining an expected time T_i the program will spend executing a method "m"
 4 if said method is not recompiled; ,
 5 determining a cost C_j of recompiling said method at an optimization level "j", for
 6 $i \leq j \leq N$; and,
 7 determining an expected time T_j the program will spend executing said method in
 8 the future, if said method is recompiled at level "j"; and,

evaluating the expression $C_j + T_j < T_i$, and, one of: recompiling "m" at level "j" if said expression is true, and, not recompiling if said expression is false.

32. The method as claimed in Claim 19, wherein said optimizing step further comprises the step of performing online feedback-directed optimizations based on feedback from the current executing program.

33. The method as claimed in Claim 32, wherein said raw profile data samples relate to call context information associated with methods called by said program, said feedback comprising said call context information.

34. The method as claimed in Claim 32, wherein said raw profile data samples relate to current program variable values, said feedback comprising a subset of values assigned to said variables during program execution.

35. The method as claimed in Claim 32, wherein said raw profile data samples relate to control flow execution within a method, said feedback comprising execution frequency of control flow paths within said executing method.

36. The method as claimed in Claim 19, wherein said execution environment includes an interpreter device.

37. A computer program product comprising a computer readable medium having recorded thereon a computer program which, when loaded in a computer, configures a computer for adaptively optimizing a computer program executing in an execution environment, said execution environment comprising one or more compiler devices for providing various levels of program optimization, said computer program executing method steps comprising:

a) sampling said executing computer program to obtain raw profile data samples;

- 7 b) characterizing said raw profile data as meeting a threshold criteria;
8 c) analyzing said characterized raw profile data for determining whether a level of
9 program optimization for said executing program is to be performed by a compiler device, and
10 generating a compilation plan in accordance with a determined level of optimization; and,
11 d) when optimization is to be performed, invoking a compiler device for
12 optimizing said executing program in accordance with said compilation plan.

1 38. The computer program product as claimed in Claim 37, wherein said characterizing step b)
2 comprises: processing said raw profile data to determine whether said data meets an activity
3 hotness threshold.

4 39. The computer program product as claimed in Claim 37, wherein said sampling step includes
5 the steps of:

6 counting said samples that are taken from the executing program; and,
7 comparing the amount of samples taken to a predetermined sampling size
8 threshold, whereby in response to an amount of samples exceeding said sampling threshold,
9 performing said characterizing step.

10 40. The computer program product as claimed in Claim 39, wherein said raw profile data
11 sampled relates to one or more method activations in said executing program.

12 41. The computer program product as claimed in Claim 40, wherein said step of processing said
1 raw profile data comprises the steps of:

2 comparing said raw profile data of method activations against said corresponding
3 activity hotness threshold for one or more methods; and
4 identifying one or more methods as meeting said activity hotness threshold.
5

1 42. The computer program product as claimed in Claim 39, further including the step of
2 adaptively adjusting said sampling size threshold.

1 43. The computer program product as claimed in Claim 40, wherein said step of optimizing
2 includes recompiling an executing method meeting an activity hotness threshold, said method
3 further including the step of adapting said sampling size threshold in accordance with an amount
4 of recompilation that occurs.

1 44. The computer program product as claimed in Claim 40, further including the step of
2 dynamically adjusting said activity hotness threshold for adapting to a current behavior of the
3 executing computer program.

1 45. The computer program product as claimed in Claim 44, wherein said step of dynamically
2 adjusting said activity hotness threshold further includes the steps of:
3 determining an amount of methods characterized as meeting said threshold criteria
4 after one or more sampling periods;
5 comparing said amount to a limit; and,
6 in response to said comparing, one of: decreasing said activity hotness threshold if
7 said limit is not met, and increasing the threshold if said limit is met or exceeded in said one or
8 more sampling periods.

1 46. The computer program product as claimed in Claim 40, further including the step of inserting
2 intrusive profiling for one or more identified methods.

1 47. The computer program product as claimed in Claim 42, wherein said step of generating a
2 compilation plan includes: providing an identifier of said method to be optimized; and, an
3 optimization level indicating a degree of optimization to be applied for said identified method.

48. The computer program product as claimed in Claim 40, wherein said analyzing step c) further includes identifying a recompilation level that minimizes expected future running time of a recompiled version.

49. The computer program product as claimed in Claim 48, wherein said step of identifying a recompilation level includes the steps of:

determining an expected time T_i the program will spend executing a method "m" if said method is not recompiled;

determining a cost C_j of recompiling said method at an optimization level "j", for $i \leq j \leq N$; and,

determining an expected time T_j the program will spend executing said method in the future, if said method is recompiled at level "j"; and,

evaluating the expression $C_j + T_j < T_i$, and, one of: recompiling "m" at level "j" if said expression is true, and, not recompiling if said expression is false.

50. The computer program product as claimed in Claim 37, wherein said optimizing step further comprises the step of performing online feedback-directed optimizations based on feedback from the current executing program.

51. The computer program product as claimed in Claim 50, wherein said raw profile data samples relate to call context information associated with methods called by said program, said feedback comprising said call context information

52. The computer program product as claimed in Claim 50, wherein said raw profile data samples relate to current program variable values, said feedback comprising a subset of values assigned to said variables during program execution.

1 53. The computer program product as claimed in Claim 50, wherein said raw profile data
2 samples relate to control flow execution within a method, said feedback comprising execution
3 frequency of control flow paths within said executing method.

1 54. The computer program product as claimed in Claim 37, wherein said execution environment
2 includes an interpreter device.

0070316-10400